

# Structured IB: Improving Information Bottleneck with Structured Feature Learning

Hanzhe Yang, Youlong Wu\*, Dingzhu We, Yong Zhou, Yuanming Shi\*,

School of Information Science and Technology, ShanghaiTech University, Shanghai, China  
 {yangzh2022, wuy11, wendzh, zhouyong, shiyu} @shanghaitech.edu.cn

## Abstract

The Information Bottleneck (IB) principle has emerged as a promising approach for enhancing the generalization, robustness, and interpretability of deep neural networks, demonstrating efficacy across image segmentation, document clustering, and semantic communication. Among IB implementations, the IB Lagrangian method, employing Lagrangian multipliers, is widely adopted. While numerous methods for the optimizations of IB Lagrangian based on variational bounds and neural estimators are feasible, their performance is highly dependent on the quality of their design, which is inherently prone to errors. To address this limitation, we introduce Structured IB, a framework for investigating potential structured features. By incorporating auxiliary encoders to extract missing informative features, we generate more informative representations. Our experiments demonstrate superior prediction accuracy and task-relevant information preservation compared to the original IB Lagrangian method, even with reduced network size.

**Code** — <https://github.com/HanzheYang19/Structure-IB>

## Introduction

The Information Bottleneck (IB) principle, a representation learning approach rooted in information theory, was introduced by (Tishby, Pereira, and Bialek 2000). The core idea of IB is to extract a feature, denoted as  $Z$ , that captures relevant information about the target  $Y$  from the input  $X$ . Specifically, IB aims to maximize the mutual information (MI) between  $Z$  and  $Y$  while constraining the MI between  $X$  and  $Z$ , effectively preserving only information essential for predicting  $Y$ . Mathematically, this is formulated as:

$$\max_{p(Z|X)} I(Z, Y) \quad \text{s.t.} \quad I(X, Z) < r, \quad (1)$$

where  $r$  controls the level of compression.

Solving this constrained optimization problem directly is challenging. To address this, IB Lagrangian methods (see, (Tishby, Pereira, and Bialek 2000; Yu et al. 2024; Rodríguez Gálvez, Thobaben, and Skoglund 2020; Shamir,

Sabato, and Tishby 2010)) introduce a positive Lagrange multiplier  $\beta$ , transforming the problem into:

$$\min_{p(Z|X)} -I(Z, Y) + \beta I(X, Z), \quad (2)$$

where  $\beta$  is typically within the range  $[0, 1]$  (Rodríguez Gálvez, Thobaben, and Skoglund 2020).

IB and its variations (Alemi et al. 2016; Kolchinsky, Tracey, and Wolpert 2019; Rodríguez Gálvez, Thobaben, and Skoglund 2020) have been applied to various domains, including image segmentation (Bardera et al. 2009), document clustering (Slonim and Tishby 2000), and semantic communication (Yang et al. 2023; Xie et al. 2023). Recent research has linked IB to Deep Neural Networks (DNNs) in supervised learning, where  $X$  represents input features,  $Y$  is the target (e.g., class labels), and  $Z$  corresponds to intermediate latent representations. Studies (Shwartz-Ziv and Tishby 2017; Lorenzen, Igel, and Nielsen 2021) have shown that IB can explain certain DNN training dynamics. Additionally, IB has been shown to improve DNN generalization (Wang et al. 2021) and adversarial robustness (Alemi et al. 2016).

Despite their practical success, IB methods face computational challenges during optimization. Variational bounds and neural estimators (Belghazi et al. 2018b; Pichler et al. 2022) have been proposed to mitigate these issues (Alemi et al. 2016; Kolchinsky, Tracey, and Wolpert 2019; Rodríguez Gálvez, Thobaben, and Skoglund 2020). However, the approximations inherent in these methods can deviate from the original objective due to the nature of upper or lower bounds and estimators.

This paper introduces the Structured Information Bottleneck (SIB) framework for IB Lagrangian methods. Unlike traditional IB, which uses a single feature, SIB explores structured features. We divide the feature extractor into a main encoder and multiple auxiliary encoders. The main encoder is trained using the IB Lagrangian, while auxiliary encoders aim to capture information missed by the main encoder. These features are combined to form the final feature. Our experiments demonstrate that SIB achieves higher  $I(Z, Y)$  for the same compression level  $I(X, Z)$ , improves prediction accuracy, and is more parameter-efficient. We also analyze how auxiliary features enhance the main feature.

\*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The contributions of our work can be summarized as follows:

- **Proposed a novel SIB framework for IB Lagrangian methods**, departing from the traditional single feature approach by incorporating multiple auxiliary encoders.
- **Developed a novel training methodology for SIB**, involving a two-stage process: initial training of the main encoder using the IB Lagrangian, followed by training auxiliary encoders to capture complementary information while maintaining distinctiveness from the main feature.
- **Demonstrated superior performance of SIB** compared to existing methods in terms of achieving higher mutual information between the compressed representation and the target variable for the same level of compression, as well as improved prediction accuracy and parameter efficiency.
- **Provided a comprehensive analysis of how auxiliary features enhance the main feature**, shedding light on the underlying mechanisms of the proposed framework.

### Relative Work

The IB Lagrangian has been extensively studied in representation learning. Several methods have been proposed to optimize it using DNNs. (Alemi et al. 2016) introduced Variational Information Bottleneck (VIB), deriving variational approximations for the IB Lagrangian:

$$I(X, Z) \leq \mathbb{E}_{q(z|x)} \log q(z|x) - \mathbb{E}_{q(t)} \log v(t), \quad (3)$$

$$I(Z, Y) \geq \mathbb{E}_{p(y,z)} \log q(y|z) + H(Y), \quad (4)$$

where  $q(\cdot|\cdot)$  and  $p(\cdot|\cdot)$  represent the probabilistic mapping and variational probabilistic mapping respectively,  $q(\cdot)$  is the probabilistic distribution,  $H(\cdot)$  is the entropy, and  $v(\cdot)$  is some prior distribution.

Departing from VIB, Nonlinear Information Bottleneck (NIB) by Kolchinsky, Tracey, and Wolpert (2019) employs kernel density estimation (Kolchinsky and Tracey 2017) to bound  $I(X, Z)$ :

$$I(X, Z) \leq -\frac{1}{N} \sum_{i=1}^N \log \frac{1}{N} \sum_{j=1}^N e^{-D_{KL}[q(z|x_i)||q(z|x_j)]}, \quad (5)$$

where  $D_{KL}(\cdot||\cdot)$  is the KL divergence and  $N$  is the total number of samples in the dataset.

However, Rodríguez Galvez (2019) demonstrated that optimizing the IB Lagrangian for different values of  $\beta$  cannot explore the IB curve when  $Y$  is a deterministic function of  $X$ . To address this, they proposed the Square IB (sqIB) by simply using the square of compression term  $(I(X, Z))^2$  instead of  $I(X, Z)$  in the Lagrangian function. Furthermore, Rodríguez Gálvez, Thobaben, and Skoglund (2020) showed that applying any monotonically increasing and strictly convex functions on  $I(X, Z)$  is able to explore the IB curve. Additionally, the authors of Yu et al. (2024) extended the IB for regression tasks with Cauchy-Schwarz divergence.

There are also other approaches that focus on neural estimators for directly approximating MI, entropy, and differential entropy (see, (Belghazi et al. 2018b; Pichler et al. 2022)),

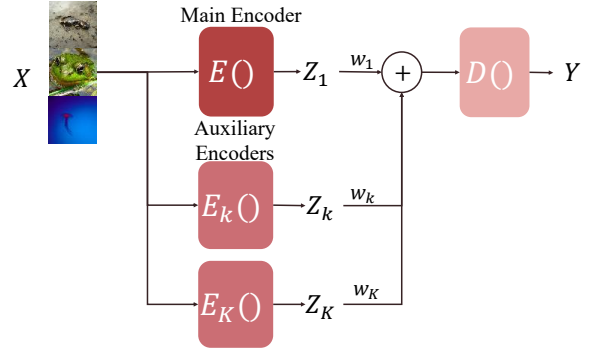


Figure 1: The illustration of the network architecture.

enabling direct IB Lagrangian optimization without variational bounds.

While these methods offer theoretical advancements, their implementations can be error-prone due to specific implementation details.

Our proposed SIB framework introduces auxiliary encoders to capture complementary information and a novel training process. Unlike prior work, we leverage the IB Lagrangian for both main and auxiliary encoder training and include an additional term to prevent feature overlap.

### Methodology

In this section, we first present the architecture of the proposed network. Subsequently, the training process will be illustrated in detail.

#### Network Architecture

Our network architecture comprises three key components: the main encoder, auxiliary encoders, and the decoder, as illustrated in Figure 1. The main encoder extracts the primary feature from the input, while auxiliary encoders capture complementary features to enrich the information content of the primary feature. These features are combined using a weighted aggregation function  $f(\cdot)$  to produce a comprehensive feature representation. In this work, we employ a simple weighted summation, which will be justified in the following subsection. The decoder processes the aggregated feature to generate the final output.

Let  $X$  and  $Y$  represent the input and label, respectively, with joint distribution  $p(X, Y)$ . The main encoder  $E$  extracts the primary feature  $Z = E(X)$ , while auxiliary encoders  $E_i$ , where  $i = 1, 2, \dots, K$  and  $K$  is the number of auxiliary encoders, extract supplementary features  $Z_i = E_i(X)$ . These features are weighted by  $w_i, i = 0, 1, \dots, K$  and aggregated as follows:

$$\hat{Z} = w_0 Z + \sum_{i=1}^K w_i Z_i, \quad (6)$$

where  $\hat{Z}$  denotes the aggregated feature. Ideally, the main feature should carry the most weight ( $w_0$  is largest). However, we do not enforce this constraint explicitly, as our experiments show that  $w_0$  naturally converges to the dominant

weight during training. The decoder  $D$  takes the aggregated feature  $\hat{Z}$  as input and produces the output  $\hat{Y} = D(\hat{Z})$ .

### Training Process

The training process is divided into three stages: training the main encoder and decoder, training auxiliary encoders, and optimizing weights.

**The main encoder and the decoder:** The training of the main encoder and decoder follows traditional IB methods like VIB (Alemi et al. 2016) or NIB (Kolchinsky, Tracey, and Wolpert 2019). The main encoder  $E$  processes the input  $X$  to generate the primary feature  $Z$ , while the decoder  $D$  directly receives  $Z$  to produce a preliminary output  $Y_0$ . This forms a Markov chain  $Z \leftrightarrow X \leftrightarrow Y_0$  (Witsenhausen and Wyner 1975; Gilad-Bachrach, Navot, and Tishby 2003). The goal is to optimize  $E$  to learn the optimal probabilistic mapping  $p(Z|X)$  and  $D$  to approximate the conditional distribution  $p(Y|Z)$  accurately. This is achieved by maximizing the MI  $I(Z, Y)$  while constraining  $I(X, Z)$  through minimizing the IB Lagrangian (Gilad-Bachrach, Navot, and Tishby 2003; Rodríguez Gálvez, Thobaben, and Skoglund 2020; Yu et al. 2024):

$$\min_{E, D} L[p(Z|X), p(Y|Z); \beta] = -I(Z, Y) + \beta I(X, Z), \quad (7)$$

where  $\beta$  denotes the Lagrange multiplier that balances the trade-off.

**The auxiliary encoders:** Auxiliary encoders are trained sequentially. For the  $i$ -th encoder  $E_i$  ( $i \in [1, K]$ ), the objective is to extract features that are informative about the target  $Y$  but independent of previously extracted features. To capture informative content, we employ the IB Lagrangian:

$$\min_{E_i} L[p(Z_i|X), p(Y|Z_i); \beta] = -I(Z_i, Y) + \beta I(X, Z_i). \quad (8)$$

Notably, to ensure compatibility and seamless integration of features extracted by different encoders, we maintain a consistent feature space throughout the network. This is achieved by fixing the decoder’s parameters after its initial training with the main encoder. Consequently, the decoder remains unchanged during the training of auxiliary encoders.

To ensure feature independence, we minimize the MI between the current feature  $Z_i$  and the concatenation of previous features  $Z + \sum_{j=1}^{i-1} Z_j$ :

$$\min_{E_i} I(Z_i, Z + \sum_{j=1}^{i-1} Z_j). \quad (9)$$

Directly minimizing the mutual information  $I(Z_i, Z + \sum_{j=1}^{i-1} Z_j)$ , equivalent to the KL divergence  $D_{KL}[p(Z_i, Z + \sum_{j=1}^{i-1} Z_j) | p(Z_i)p(Z + \sum_{j=1}^{i-1} Z_j)]$ , is computationally challenging due to the complex distributions involved, each comprising a mixture of numerous components (Pan et al. 2021). To address this, we adopt a sampling-based approach.

First, we generate samples from the joint distribution  $p(Z_i, Z + \sum_{j=1}^{i-1} Z_j)$  by randomly selecting inputs  $X$  from

the dataset and extracting the corresponding features. Subsequently, we obtain samples from the product distribution  $p(Z_i)p(Z + \sum_{j=1}^{i-1} Z_j)$  by shuffling the samples from the joint distribution along the batch dimension (Belghazi et al. 2018a).

To estimate the KL divergence, we employ the density-ratio trick (Nguyen, Wainwright, and Jordan 2007; Kim and Mnih 2018). A discriminator  $d$  is introduced to distinguish between samples from the joint and product distributions. The discriminator is trained adversarially using the following objective:

$$\min_{E_i} \max_d \mathbb{E}_{p(Z_i)p(Z + \sum_{j=1}^{i-1} Z_j)} \log d(Z_i, Z + \sum_{j=1}^{i-1} Z_j) + \mathbb{E}_{p(Z_i, Z + \sum_{j=1}^{i-1} Z_j)} \log(1 - d(Z_i, Z + \sum_{j=1}^{i-1} Z_j)). \quad (10)$$

It should be noted that the  $I(Z_i, Z + \sum_{j=1}^{i-1} Z_j)$  is minimized when the Nash equilibrium is achieved (Goodfellow et al. (2014)).

**The weights:** The weights  $w_i$  are optimized to fine-tune the overall network using the IB Lagrangian. The objective is to minimize:

$$\min_{\{w_i\}_{i \in [0, K]}} L[p(\hat{Z}|X), p(Y|\hat{Z}); \beta] = -I(\hat{Z}, Y) + \beta I(X, \hat{Z}). \quad (11)$$

It should be noted that the reason why we use weights to be the fine-tune parameter instead of retraining the decoder is that, retraining the decoder could potentially alter the learned relationships between features and the target variable  $Y$ , affecting the calculated MI  $I(Z_i, Y)$ . By using weights, we preserve the integrity of the trained features while optimizing their combined representation.

Detailed implementation specifics are provided in the supplementary materials.

### Justification of $f(\cdot)$

Any function that increases  $I(\hat{Z}, Y)$ , where  $\hat{Z} = f(Z, [Z_i])$ , but does not necessarily increase  $I(X, \hat{Z})$ , will be appropriate. Here we only justify the choice of weighted sum.

**Theorem 1** Assume that  $Z, Z' \in \mathbb{R}^D$  are independent, where  $Z \sim \mathcal{N}(\mu, \Sigma)$ ,  $Z' \sim \mathcal{N}(\mu', \Sigma')$ ,  $D$  is the dimension,  $\mu, \mu' \in \mathbb{R}^D$  are the means, and  $\Sigma, \Sigma' \in \mathbb{R}^{D \times D}$  are the diagonal positive definite covariance matrices. Moreover, let  $d(Z) = WZ : \mathbb{R}^D \rightarrow \mathbb{R}^D$  be the linear decoder function with full-rank parameter  $W \in \mathbb{R}^{D \times D}$ ,  $h(\cdot)$  is the one-hot coding function, and  $Y' = d(Z + Z')$ . Then, given data  $X$  and its target  $Y$ , we have

$$I(Z' + Z, Y) \geq I(Z, Y), \quad (12)$$

when the following conditions are satisfied:

$$I(h(Y'), Y) = H(Y), \quad (13)$$

$$\det(\Sigma') \geq \frac{1}{(2\pi e)^D}. \quad (14)$$

Theorem 1 posits that auxiliary features can augment the information content of the original feature regarding  $Y$  by considering  $wZ + \sum_{j=1}^{i-1} w_j Z_j$  as  $Z$  and  $w_i Z_i$  as  $Z'$  within the SIB framework.

To validate the assumptions underlying Theorem 1, we employ variational encoders as suggested by Alemi et al. (2016) to satisfy the Gaussianity assumption. The independence condition is enforced through the minimization of  $I(Z_i, Z + \sum_{j=1}^{i-1} Z_j)$ . A one-layer MLP decoder with the dimension of  $Z$  matching the number of classes fulfills the decoder function assumption. Achieving global optimality during training ensures that  $I(h(Y'), Y) = H(Y)$  (Pan et al. 2021, Proof of Theorem 2). The condition  $\det(\Sigma') \geq \frac{1}{(2\pi e)^D}$  typically holds true due to the large number of classes in practical scenarios. Furthermore, when using NIB, the covariance matrix  $\Sigma$  can be treated as a hyperparameter rather than a trainable parameter.

Regarding  $I(X, \hat{Z})$ , we will empirically demonstrate that incorporating auxiliary features remains at a similar level in terms of  $I(X, \hat{Z})$  compared to  $I(X, Z)$ . More effort on the theoretical analysis of  $I(X, \hat{Z})$  will be put on.

We also provide an intuitive explanation of the weighted summation mechanism. Figure 2 visually represents the feature spaces of the original IB Lagrangian method and our SIB approach. In the standard IB Lagrangian method, the encoder generates a single feature,  $Z$ , confined to a specific subspace (leftmost circle). While theoretically capable of capturing sufficient information, practical implementations and approximation bounds might lead to information loss. Our SIB framework introduces an auxiliary feature,  $Z'$ , expanding the feature space (middle circle). Initially, the subspaces spanned by  $Z$  and  $Z'$  may overlap significantly. Through the proposed training stages, these subspaces become more distinct as we encourage feature independence *i.e.*,  $(I(Z, Z') < \epsilon$  for a small constant  $\epsilon$ ). The subsequent weight optimization further refines the combined subspace spanned by  $Z$  and  $Z'$  to maximize information about  $Y$  as shown in the rightmost circle.

While the weighted aggregation employed in this work effectively combines features, it is acknowledged that a more sophisticated operator capable of fully capturing the intricacies of the feature space could potentially yield even more refined and accurate representations. Future research will explore the development of such operators.

## Discussion

While primarily developed to augment IB Lagrangian methods, our proposed framework is theoretically applicable to any method satisfying the conditions outlined in Theorem 1. Adapting this framework to other methods would primarily involve modifying the encoders and incorporating the independence-encouraged component into the respective loss function. However, due to the prevalence of IB Lagrangian approaches and computational constraints, our current research focuses on IB Lagrangian implementation.

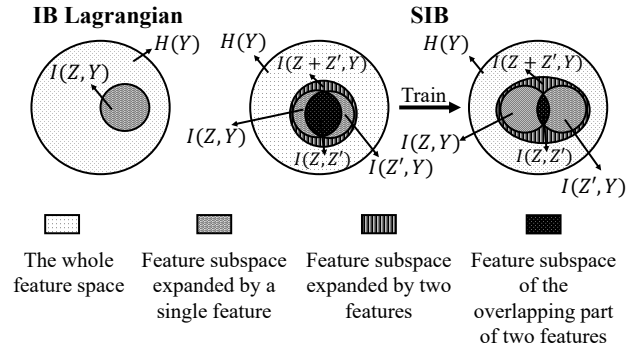


Figure 2: Illustration of the feature space. **Left:** The entire feature space (large circle) containing the feature subspace (ellipse) generated by a single feature vector using the IB Lagrangian method. **Middle:** The feature subspace (small circle) spanned by two vectors from an untrained SIB within the overall feature space (large circle). Due to substantial overlap, this subspace is limited in its coverage. **Right:** The feature subspace (small circle) is spanned by two vectors from a well-trained SIB within the overall feature space (large circle). Here, the feature subspace is expanded by two minimally overlapping features, maximizing its information content, as demonstrated in Theorem 1.

## Experiments

In this section, we will demonstrate our experimental setups and several results. Specifically, we first evaluate and analyze the performance of the proposed SIB in comparison with its corresponding IB algorithm. Next, we examine the behavior of weights in the SIB, followed by an analysis of both SIB and corresponding IB methods on the IB plane. Due to space limitations, experiments on encoder dropout, which show the significance of SIB weights, are provided in the Appendix.

### Experiment Setups

**Dataset:** Several benchmark datasets are used: MNIST (LeCun et al. (1998)) and CIFAR10 (Krizhevsky, Hinton et al. (2009)). All the results are based on the test set.

**Comparing algorithms:** Existing variants of the IB Lagrangian, such as VIB, square VIB (sqVIB), and NIB, are compared to the proposed structured counterparts, SVIB, sqSVIB, and SNIB. Additionally, the estimators MINE and KNIFE, along with their structured versions, are evaluated, whose detailed comparisons is provided in the Appendix.

**Network structure:** PyTorch (Paszke et al. 2019) is employed to implement the algorithms. Network architectures vary based on the specific algorithm and dataset. In particular, a multilayer perceptron (MLP) is used for MNIST, while a convolutional neural network is adopted for CIFAR-10. Detailed network architectures and hyperparameter settings are provided in the Appendix.

**Devices:** The experiments were conducted on a system equipped with a 12th Generation Intel(R) Core(TM) i7-12700 CPU, operating at a frequency of 2100MHz. The

CPU has 12 cores and 20 logical processors. The system also includes an NVIDIA GeForce RTX 4070 Ti GPU with a WDDM driver and a total available GPU memory of 12282MB.

### Performance Analysis and Evaluation

The accuracy,  $I(Z, Y)$ ,  $I(X, Z)$ , and the number of model parameters with respect to the number of encoders are depicted in Figure 3. The Lagrange multiplier  $\beta$  is set to 1 for SVIB and sqVIB, and 0.01 for SNIB. Notably, when  $K = 1$ , the algorithms reduce to the standard IB Lagrangian. Furthermore, to estimate  $I(X, Z)$ , we employ Monte Carlo sampling (Goldfeld et al. (2018)). Besides, given that  $I(Z, Y) = H(Y) - H(Y|Z)$ , we estimate the conditional entropy  $H(Y|Z)$  using the cross-entropy loss. It is also important to note that  $H(Y)$  is a known constant determined by the dataset. Indeed, as depicted in Figure 3, the upper four diagrams represent results based on the MNIST dataset, while the lower four diagrams are derived from the CIFAR10 dataset.

Figure 3 illustrates the accuracy,  $I(Z, Y)$ ,  $I(X, Z)$ , and model parameter count in relation to the number of encoders. The Lagrange multiplier,  $\beta$ , is set to 1 for SVIB and sqVIB, and 0.01 for SNIB. Importantly, the algorithms reduce to the standard IB Lagrangian when  $K = 1$ . MI  $I(X, Z)$  is estimated using Monte Carlo sampling (Goldfeld et al. 2018), while  $I(Z, Y) = H(Y) - H(Y|Z)$  is derived from the conditional entropy  $H(Y|Z)$ , calculated via cross-entropy loss and a dataset-dependent constant  $H(Y)$ . The upper and lower rows of Figure 3 present results for MNIST and CIFAR-10, respectively.

For MNIST dataset, our structured algorithms consistently outperform the original IB Lagrangian in terms of accuracy and  $I(Z, Y)$  even when the number of model parameters is reduced. This improvement is attributed to our method’s ability to extract supplementary features that enhance the information content of the main feature. However, as  $K$  increases, the performance may meet the bottleneck which will prevent it from surging. Thus, the number of encoders which determine the amount of parameters should be chosen carefully. In most cases,  $I(X, Z)$  experiences a slight decrease as well. However, the mechanism behind this phenomenon should be investigated in the future. The performance on the CIFAR10 dataset is similar to that of MNIST.

For the MNIST dataset, our structured algorithms consistently surpass the original IB Lagrangian in terms of accuracy and  $I(Z, Y)$ , even with reduced model parameters. This enhancement is attributed to our method’s capacity to extract complementary features that enrich the primary feature’s information content. Nevertheless, increasing the number of encoders  $K$  may lead to performance plateaus, necessitating careful selection of this hyperparameter. While  $I(X, Z)$  generally exhibits a slight decline, the underlying cause of this behavior warrants further exploration. Notably, the CIFAR-10 dataset demonstrates performance trends similar to MNIST.

In summary, our proposed method consistently outperforms traditional approaches in terms of accuracy and feature extraction, even when using smaller network archi-

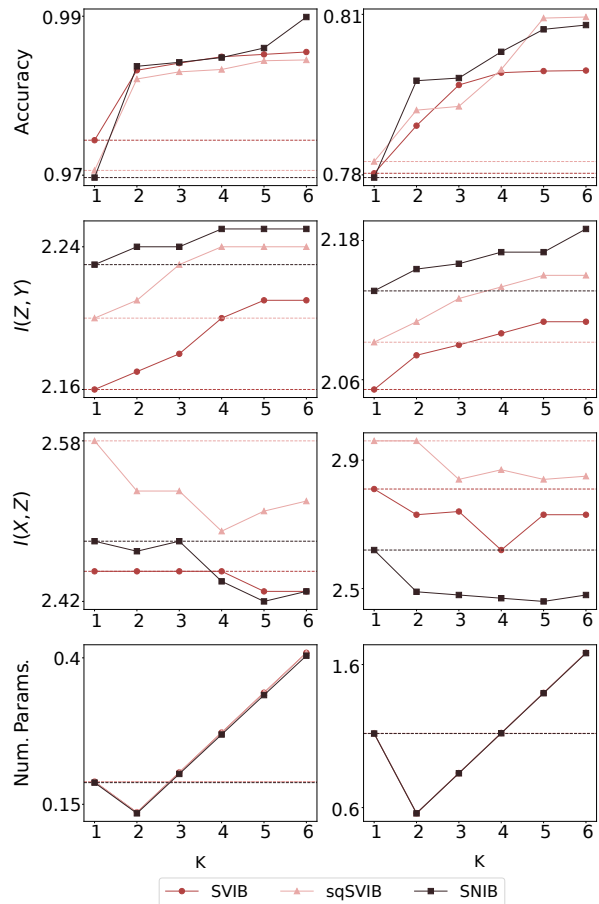


Figure 3: The accuracy, MI  $I(Z, Y)$ ,  $I(X, Z)$  and the numbers of model parameters (Num. Params., in Millions) v.s. the number of encoders. The IB Lagrangian versions are marked by dashed lines. The left four figures are based on MNIST while the figures in the right are the results of CIFAR10.

tures. To optimize performance and computational efficiency, careful selection of the hyperparameter  $K$  is crucial, as diminishing returns may occur with excessively large values.

### Behavior of Weights

In this experiment, the Lagrange multiplier  $\beta$  was set to 1 for SVIB and sqVIB, and 0.01 for SNIB. Additionally, the achieved accuracies consistently matched those reported in the previous subsection.

Intriguingly, we observed a consistent pattern where the main encoder was assigned the largest weight, even without explicit constraints. As detailed in Table 1, the weights collectively approximated 1 despite lacking explicit normalization. Furthermore, normalizing the weights to sum exactly to one had negligible impact on performance metrics. Consequently, additional weight manipulations were deemed unnecessary.

MNIST			
K	SVIB	sqSVIB	NIB
2	(0.9453, 0.0833)	(0.9302, 0.1041)	(0.9498, 0.0907)
3	(0.8901, 0.0441, 0.1159)	(0.9202, 0.0407, 0.0629)	(0.7927, 0.0729, 0.1438)
CIFAR10			
	SVIB	sqSVIB	NIB
2	(0.8061, 0.2383)	(0.9345, 0.0872)	(0.8323, 0.2574)
3	(0.8037, 0.0889, 0.1659)	(0.6684, 0.2058, 0.2626)	(0.8457, 0.2040, 0.1275)

Table 1: The weights of encoders.

Furthermore, as detailed in Table 1, the auxiliary encoder weights are generally larger for the CIFAR-10 dataset compared to MNIST, indicating that these encoders contribute more significantly to the overall performance on more complex tasks.

### Behavior on IB Plane

This section analyzes the behavior of our proposed algorithms and original IB Lagrangian algorithms on the IB plane, using two encoders, i.e.,  $K = 2$ . Figure 4 presents results for MNIST (left column) and CIFAR-10 (right column) datasets. Our methods consistently achieve higher  $I(Z, Y)$  compared to the original IB Lagrangian. However, the rate of decrease in  $I(X, Z)$  is generally slower. Notably, SNIB consistently exhibits lower  $I(X, Z)$  than NIB. Furthermore, the  $I(X, Z)$  of SVIB and sqSVIB converges to that of VIB and sqVIB as the Lagrange multiplier  $\beta$  increases. Crucially, our methods demonstrate the ability to attain higher  $I(Z, Y)$  for a given level of  $I(X, Z)$  as revealed in the third row.

In summary, our proposed method effectively improves the informativeness of extracted features while maintaining or even enhancing data compression (especially for SNIB). Additionally, our models achieve superior feature utility at comparable compression levels compared to the original IB Lagrangian. Significantly, these advantages are realized with a reduced number of model parameters as shown in Figure 3.

### Conclusion and Discussion

This paper introduces a novel structured architecture designed to improve the performance of the IB Lagrangian method. By incorporating auxiliary encoders to capture additional informative features and combining them with the primary encoder’s output, we achieve superior accuracy and a more favorable IB trade-off compared to the original method, even with a significantly smaller network. Our results highlight the effectiveness of our approach in enhancing feature utility while maintaining desired compression levels. Furthermore, we demonstrate the crucial role of auxiliary encoders in driving these improvements.

Several limitations warrant further investigation. The current linear aggregation function may oversimplify feature space interactions, necessitating more sophisticated combination strategies. Additionally, theoretical advancements,




Figure 4: Behavior of Algorithms on the IB Plane. The original IB Lagrangian methods and their structured counterparts are represented in the same color, differentiated by solid and dashed lines. The left and right figures correspond to the MNIST and CIFAR10 datasets, respectively.

such as relaxing the Gaussian assumption and investigating how to assist the compression, could refine both the aggregation function and decoder design. While this study focuses on IB Lagrangian methods, our structured architecture holds potential for broader applications. Overall, we believe our proposed framework offers a promising foundation for future research.

### References

- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Bardera, A.; Rigau, J.; Boada, I.; Feixas, M.; and Sbert, M. 2009. Image segmentation using information bottleneck method. *IEEE Transactions on Image Processing*, 18(7): 1601–1612.
- Belghazi, M. I.; Baratin, A.; Rajeshwar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, D. 2018a. Mutual infor-

- mation neural estimation. In *International conference on machine learning*, 531–540. PMLR.
- Belghazi, M. I.; Baratin, A.; Rajeswar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, R. D. 2018b. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*.
- Gilad-Bachrach, R.; Navot, A.; and Tishby, N. 2003. An information theoretic tradeoff between complexity and accuracy. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003. Proceedings*, 595–609. Springer.
- Goldfeld, Z.; Berg, E. v. d.; Greenewald, K.; Melnyk, I.; Nguyen, N.; Kingsbury, B.; and Polyanskiy, Y. 2018. Estimating information flow in deep neural networks. *arXiv preprint arXiv:1810.05728*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Kim, H.; and Mnih, A. 2018. Disentangling by factorising. In *International conference on machine learning*, 2649–2658. PMLR.
- Kolchinsky, A.; and Tracey, B. D. 2017. Estimating mixture entropy with pairwise distances. *Entropy*, 19(7): 361.
- Kolchinsky, A.; Tracey, B. D.; and Wolpert, D. H. 2019. Nonlinear information bottleneck. *Entropy*, 21(12): 1181.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Lorenzen, S. S.; Igel, C.; and Nielsen, M. 2021. Information bottleneck: Exact analysis of (quantized) neural networks. *arXiv preprint arXiv:2106.12912*.
- Nguyen, X.; Wainwright, M. J.; and Jordan, M. 2007. Estimating divergence functionals and the likelihood ratio by penalized convex risk minimization. *Advances in neural information processing systems*, 20.
- Pan, Z.; Niu, L.; Zhang, J.; and Zhang, L. 2021. Disentangled information bottleneck. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 9285–9293.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Pichler, G.; Colombo, P.; Boudiaf, M.; Koliander, G.; and Piantanida, P. 2022. KNIFE: Kernelized-Neural Differential Entropy Estimation.
- Rodríguez Galvez, B. 2019. The Information Bottleneck: Connections to Other Problems, Learning and Exploration of the IB Curve.
- Rodríguez Gálvez, B.; Thobaben, R.; and Skoglund, M. 2020. The convex information bottleneck lagrangian. *Entropy*, 22(1): 98.
- Shamir, O.; Sabato, S.; and Tishby, N. 2010. Learning and generalization with the information bottleneck. *Theoretical Computer Science*, 411(29-30): 2696–2711.
- Shwartz-Ziv, R.; and Tishby, N. 2017. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*.
- Slonim, N.; and Tishby, N. 2000. Document clustering using word clusters via the information bottleneck method. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, 208–215.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Wang, Z.; Huang, S.-L.; Kuruoglu, E. E.; Sun, J.; Chen, X.; and Zheng, Y. 2021. Pac-bayes information bottleneck. *arXiv preprint arXiv:2109.14509*.
- Witsenhausen, H.; and Wyner, A. 1975. A conditional entropy bound for a pair of discrete random variables. *IEEE Transactions on Information Theory*, 21(5): 493–501.
- Xie, S.; Ma, S.; Ding, M.; Shi, Y.; Tang, M.; and Wu, Y. 2023. Robust information bottleneck for task-oriented communication with digital modulation. *IEEE Journal on Selected Areas in Communications*.
- Yang, Y.; Wu, Y.; Ma, S.; and Shi, Y. 2023. Multi-Task-Oriented Broadcast for Edge AI Inference via Information Bottleneck. In *GLOBECOM 2023-2023 IEEE Global Communications Conference*, 3142–3147. IEEE.
- Yu, S.; Yu, X.; Løkse, S.; Jenssen, R.; and Principe, J. C. 2024. Cauchy-Schwarz Divergence Information Bottleneck for Regression. *arXiv preprint arXiv:2404.17951*.